A SAFETY-BASED DECISION MAKING ARCHITECTURE FOR AUTONOMOUS SYSTEMS

MS

NAGW-1333

by

Joseph C. Musto and L.K. Lauderbaugh

Rensselaer Polytechnic Institute
Department of Mechanical Engineering,
Aeronautical Engineering, and Mechanics
Troy, New York 12180-3590

July 1991

CIRSSE REPORT #98

A Safety-Based Decision Making Architecture for Autonomous Systems

Joseph C. Musto and L.K. Lauderbaugh
Department of Mechanical Engineering, Aeronautical Engineering, and Mechanics
Rensselaer Polytechnic Institute
Troy, New York 12180-3590

July 31, 1991

Abstract

Engineering systems designed specifically for space applications often exhibit a high level of autonomy in the control and decision-making architecture. As the level of autonomy increases, more emphasis must be placed on assimilating the safety functions normally executed at the hardware level or by human supervisors into the control architecture of the system. This paper details the development of a decision-making structure which utilizes information on system safety. A quantitative measure of system safety, called the *safety self-information*, is defined. This measure is analogous to the reliability self-information defined by McInroy and Saridis, but includes weighting of task constraints to provide a measure of both reliability and cost. An example is presented in which the safety self-information is used as a decision criterion in a mobile robot controller. The safety self-information is shown to be consistent with the entropy-based Theory of Intelligent Machines defined by Saridis.

1 Introduction

Safe operation is a consideration whenever an engineering system is designed and constructed. Research in the safety of robotic systems has been concentrated in three areas: human factors, such as the layout of control panels, teach pendants, and mechanical guards; robot factors, such as perimeter safety zones and "watchdog" safety systems; and systems issues, such as fault-tree analysis of robot accidents and operator training [1]. Each of these issues can be categorized as "hardware level" approaches to safety; the goal of these approaches is to minimize the risk of accidents caused by human interference with the robotic system, and provide emergency shutdown of the system when an accident is imminent or has occurred.

Although each of these safety issues may be relevant in the construction of highly autonomous and fully autonomous systems, the ideas are generally drawn from safety approaches in fixed automation systems, which operate within highly specified physical constraints over a well-defined set of parameters. They fail to address the needs of highly autonomous systems, particularly those designed to perform ill-defined tasks in unstructured environments. In addition, only the safety of the human operator is considered; in autonomous systems, the safety of the system with regard to environmental hazards must also be taken into account. Consider, for example, a mobile robotic platform operating as an exploration vehicle on unknown terrain. Although it would be necessary to provide standard safety features, such as a bumper system hardwired to stop the drive motors in the event of collision, other standard safety features would fall short in fully safeguarding the system. A safety fence cannot be built around the terrain to be explored; the controller of the robotic system must be capable of assessing potential environmental hazards and making control decisions with this hazard assessment in mind. In addition, the controller must be capable of weighing potential risks to the system with the urgency of the task to be performed; the controller should be capable of making a control decision when it may become necessary to violate an operating specification in order to complete an urgent task.

This paper presents a method for assessing the level of safety of various plans for performing a task in an autonomous system. A quantity known as the *safety self-information* (SSI) will be introduced. This quantity will be a reflection of both the probability that a plan will violate a task specification,

as well as the potential hazard to the system caused by violating that specification. This work is based on the reliability analysis for Intelligent Machines formulated by McInroy and Saridis [2-4]. The approach will be demonstrated in a case study of a mobile robot performing a task with a dynamic obstacle in the environment. In addition, since the safety analysis is to be used as a decision-making tool within the Hierarchical Control structure for intelligent machines, the SSI will be shown to be consistent with the principle of Increasing Precision with Decreasing Intelligence [5,6].

2 Safety Analysis for Autonomous Systems

Safety analysis for autonomous systems is concerned with selecting a plan for executing a specified task based on minimizing the potential risk to the system. The analysis is probabilistic in nature; it is assumed that knowledge obtained from sensors and contained in the data base of the autonomous system controller contains a degree of uncertainty, and can be modeled as a random variable. Safety analysis is based on reliability theory, but provides augmentation of reliability with cost information to establish a measure of risk to the autonomous system. In this section, a review of reliability theory will be presented. From this background, a method of safety analysis for autonomous systems will be proposed. The Theory of Intelligent Machines will be introduced, and the proposed safety analysis will be shown within the structure of intelligent machines.

2.1 Reliability Analysis

In order to develop safety analysis for autonomous systems, a review of reliability theory is necessary. Safety analysis uses as its basis the following definition of structural reliability, presented by Ang and Tang [7], and applied to Intelligent Machines by McInroy and Saridis [2-4].

Consider a system whose states are defined by a set of i random variables, x_i. These states represent sensor data or knowledge contained in a data base for use by the intelligent controller of the autonomous system. The task to be performed is described by a series of performance functions, which are functions of the system states, each denoted $g(X)$. The performance functions are defined such that if the specification is not violated, then $g(X) > 0$;

failing to meet the specification results $g(X) \leq 0$. With these definitions, the reliability index β is defined as the minimum distance between the origin of a set of uncorrelated standard normal variates derived from the state variables X and the failure surface $g(X) = 0$. Physically, β can be thought of as the "distance" between the current state of the system and a state at which the specification in question would be violated. In the case where the x_i are uncorrelated Gaussian random variables with mean μ_i and standard deviation σ_i, and g(X) is a linear specification of the following form:

$$g(X) = a_0 + \sum_i a_i x_i \tag{1}$$

the reduced variates can be determined by:

$$x_i' = \frac{x_i - \mu_i}{\sigma_i}, i = 1, 2, ..., n \tag{2}$$

and the reliability index β can be determined by:

$$\beta = \frac{a_0 + \sum_i a_i \mu_i}{\sqrt{\sum_i (a_i \sigma_i)^2}} \tag{3}$$

From this, the reliability can be measured by:

$$R = \Phi(\beta) \tag{4}$$

where $\Phi(\cdot)$ is the normal cumulative distribution function. Methods for calculating the reliability index and reduced variates for other standard distributions and nonlinear specifications can be found in Ang and Tang[7].

The reliability calculated using this method is known as the *system reliability*, and can be interpreted as the probability that a given task specification will not be violated. For tasks with multiple specifications, reliabilities of parallel specifications must be combined using the following relationship:

$$R_p = 1 - \prod_i (1 - R_i) \tag{5}$$

After reducing parallel reliabilities such that only a set of series reliabilities remain, the overall reliability of a system performing the specified task can be computed as follows:

$$R_{tot} = R_1 R_2 ... R_n \tag{6}$$

McInroy and Saridis propose that for an Intelligent Machine, this concept of system reliability can be viewed as a flow of reliability information through the Hierarchical Control structure. They define the *reliability self-information* (RSI), denoted $I(R)$, as follows:

$$I(R) = -log(R) \tag{7}$$

It can be shown that by definition of RSI, reliability can be treated in a framework consistent with the Theory of Intelligent Control [2,9]. By evaluating the RSI for a list of plans generated by an autonomous system controller and selecting the most reliable plan, an intelligent control system can use this reliability analysis as a design tool [2-4].

Reliability analysis can be used to determine the probability that a task specification is met. However, in using the RSI as a means of selection of a plan for task execution, it is implied that all task specifications are of equal importance; no information regarding the priority of specifications or the cost of violating a given specification are included in the analysis. Consider the situation where multiple specifications define a given task, i.e. a robot performing a peg insertion with specifications on gripper position, gripper overshoot, and execution time. Implicitly, there are economic costs associated with the violation of constraints; in the example case, assume that violation of the position and overshoot constraints will cause damage to the workpiece, while violation of the execution time specification will result in a delay of mission and increases in mission costs. In this case, a decision based solely on RSI will ignore the costs associated with the specifications; perhaps an alternative analysis could be performed which would weigh the relative costs of workpiece replacement and mission time, prioritize the specifications based on this weighting, and calculate some decision index analogous to the RSI but including a weighting function. The following analysis will result in a quantity defined as the *safety self-information* (SSI), which can be viewed as a weighted measure of reliability. It is proposed that this SSI quantity can be used in a decision-making structure of an Intelligent Machine.

2.2 Safety Analysis and Safety Self-Information

Consider the system used in the derivation of RSI presented in Section 2.1: a system whose states are represented by n uncorrelated Gaussian random

variables x_i, each with a known expected value μ_i and standard deviation σ_i, designated to perform a task described by m specifications, $g_k(X)$. With each specification, there is an economic cost associated with violation of that specification, denoted C_k. Using these definitions, a safety analysis resulting in the calculation of the SSI will be derived. From the standpoint of an autonomous system, safety analysis will be defined as the measurement and reduction of risk to an autonomous system. In the course of the analysis, risk is measured as the penalty incurred by the system when a specification is violated. Often, this penalty is defined as an economic cost, such as the cost of replacing a part damaged when a specification is violated, or the cost of repeating a task which is performed improperly. In the presentation of this analysis, this economic definition of penalties will be used; however, it should be noted that cost information is used as a relative weighting function, and non-economically based weighting functions may be substituted for economic cost information in the analysis.

The philosophy of the safety analysis is as follows: in the calculation of the reliability of a plan, the statistics of the random variables describing system states are used to calculate the probability that a given constraint will not be violated. In order to focus the analysis on constraints which are most costly, information regarding the states of the system will be treated as more uncertain when misestimated state information could result in greater risk to the system. To accomplish this, the standard deviations of the state variables are modified according to the weighting of the constraint being analyzed; the standard deviations of the state variables are increased proportionally to increased cost. This has the effect of "stretching out" the distributions of state variables when calculating the probability of violating costly constraints; in essence, risk is introduced into the safety analysis by assuring that costly constraints are met with a greater "factor of safety". By introducing a higher level of uncertainty into the analysis in areas of greater risk, reliability information is augmented with cost information.

The calculation of the SSI is as follows: numerical cost values for each constraint are normalized to provide a measure of relative costs. These relative costs are defined by:

$$c_k = \frac{C_k}{C_{\min}} \tag{8}$$

where $C_{\min}$ is the minimum cost of all C_k. This relative cost value is then used to modify the distribution of all state variables used in the specification

$g_k(X)$. It is used as a multiplier for the standard deviation, yielding a term analogous to the reliability index, known as the *safety index* Ψ, computed as follows:

$$\Psi = \frac{a_0 + \sum_i a_i \mu_i}{\sqrt{\sum_i (a_i \sigma_i c_k)^2}} \tag{9}$$

Utilizing the normalized, zero-mean, Gaussian cumulative distribution function, the *safety factor*, S, can be computed:

$$S = \Phi(\beta) \tag{10}$$

Similarly to the RSI, the SSI, denoted $F(S)$ is computed as follows:

$$F(S) = -log(S) \tag{11}$$

Physically, the SSI can be viewed as a measure providing a more conservative estimate of system reliability, in which specifications carrying a greater risk are met with a higher degree of certainty. Numerically, the SSI can be used as an index on which to base safety-related decisions in the control structure of an intelligent machine. An illustrative example is provided in Section 3.

2.3 Safety Self-Information and the Theory of Intelligent Machines

The safety analysis presented in this paper is intended to be used as a design and analysis tool for the control of autonomous systems. The development of general tools for the design of Intelligent Machines has been addressed by Saridis [5,6]. The method proposed by Saridis is summarized in the *Theory of Intelligent Machines.* In this section, it will be shown that safety analysis based on the principle of the SSI is consistent with the general framework of the Theory of Intelligent Machines, and can be integrated into hierarchical control structures developed utilizing the principles of this theory.

The Theory of Intelligent Machines is a design and analysis method developed by Saridis to provide a theoretical structure for intelligent control systems. The theory unifies concepts from Artificial Intelligence, Operations Research, and Control Theory; in this theory, machine intelligence is modeled as a flow of information through the hierarchical control structure of the Intelligent Machine [5,6]. A fundamental concept of the Theory of Intelligent Machines is the *Principle of Increasing Precision with Decreasing*

Intelligence. It will be shown in this section that the information provided by the safety self-information quantity is consistent with this principle, and can be used within the hierarchical structure of the Intelligent Machine.

In short, the Principle of Increasing Precision with Decreasing Intelligence states that Machine Intelligence (MI) operates on facts in a database (DB) to produce a rate of knowledge flow in the machine (R):

$$(MI) : (DB) \Rightarrow (R)$$

This implies that for a constant rate of knowledge R, machine intelligence is larger for a small database. As shown by McInroy and Saridis [2], reliability self-information can be interpreted within the framework of this principle; at the low levels of an intelligent machine, a decrease in the size or accuracy of the database must be countered with an increase in control performance to maintain a constant RSI. The same can be said to be true for the SSI. As shown in Equations 9-11, the SSI is shown to be directly proportional to both the uncertainty of the state variable measurements and the costs associated with the task specifications; for measurements with a large variance or specifications with a large associated cost, the SSI becomes large, indicating a decreased level of safety. To counter this decrease in the level of safety caused by an increase the uncertainty of information in the database, increased control performance must be obtained. In this manner, the Principle of Increasing Precision with Decreasing Intelligence is shown to be applicable to analysis using the SSI. In addition, since the safety analysis makes use of a self-information term calculated on a logarithmic scale, it can be described by the same mathematical properties as entropy. This interpretation of SSI as an analog to entropy provides a convenient method for incorporating SSI into the information theoretic setting of the Theory of Intelligent Machines.

3 Example: A Safety-Based Decision Structure for a Mobile Robot

In this section, the safety analysis presented in Section 2 will be applied to a simplified problem which is representative of the type encountered in an autonomous mobile robotic environment. The results of a reliability analysis will be contrasted with the results of the safety analysis. The analysis will

be shown to be consistent with the structure of the Theory of Intelligent Machines.

3.1 Problem Statement

A mobile robot, r_1, is operating in an environment with a dynamic obstacle, r_2 (see Figure 1). The positions of r_1 and r_2 are known exactly, as shown in Figure 1. It is known with perfect certainty that r_2 is traveling at 3 m/s along a straight path perpendicular to the path of r_1, which is also straight. The velocity of the robot r_1 can be obtained from sensors; the sensor currently reads 5 m/s, and the sensor information is known to be normally distributed with a standard deviation of 0.1 m/s. A collision between r_1 and r_2 will result in damage to the bumper of r_1, which will yield repair costs of \$500. The mission to be completed is to transport collected soil samples out of the collection area before contamination occurs; therefore, r_1 must move at least 4.5 m along the current path in 1 s. If contamination occurs, the mission will have to be repeated and more soil samples will need to be collected, at a cost of \$200. Additionally, it is known that the drive motor of r_1 has speed limitations, and the motor will be damaged at velocities greater than 6.1 m/s. Motor replacement bears a cost of \$900. At this stage of autonomous planning, the intelligent control structure must be used to select an acceleration profile for r_1. Three options are available: accelerate at 1 m/s^2, maintain constant velocity, or decelerate at 1 m/s^2. It is assumed for simplicity that the decision will be made instantaneously, and cannot be changed again during the course of operation.

In order to proceed with reliability and safety analysis, the task specifications must be posed in standard notation. Using the format introduced in the previous section, the task specifications can be written as a set of four constraint equations:

For the velocity specification:

$$g_1(v, a) = 6.1 - v - at > 0 \tag{12}$$

For the mission specification:

$$g_2(v, a) = vt + \frac{1}{2}at^2 - 4.5 > 0 \tag{13}$$

For collision avoidance:

$$g_3(v,a) = 5 - vt - \frac{1}{2}at^2 > 0 \tag{14}$$

$$g_4(v,a) = vt + \frac{1}{2}at^2 - 5 > 0 \tag{15}$$

where v is the velocity of r_1, a is the acceleration of r_1, and t is the elapsed time. For simplicity, we will consider a time interval of 1 s. Let the three acceleration profiles (accelerate, maintain constant velocity, and decelerate) be denoted P_1, P_2, and P_3, respectively.

3.2 Reliability Analysis

Using reliability analysis, each of the three plans (P_1, P_2, and P_3) will be evaluated. The plan with the smallest RSI, I(R), will be selected as the most reliable plan. The reliability of P_1, where $a = 1m/s^2$, can be determined as follows:

For specification $g_1(a,v) = 5.1 - v > 0$:

$$\beta_1 = \frac{5.1 - \mu_v}{\sqrt{((-1)(1))^2}} \tag{16}$$

With $\mu_v = 5m/s$, this can be evaluated as:

$$\beta_1 = 1.0$$

Evaluating the cumulative distribution function:

$$R_1 = \Phi(\beta_1) \tag{17}$$

$$= 0.8413$$

Repeating this analysis for each of the remaining three specifications yields:

$$\beta_2 = 9.0; R_2 \simeq 1$$

$$\beta_3 = -5.0; R_3 \simeq 0$$

$$\beta_4 = 5.0; R_4 \simeq 1$$

Since specifications g_3 and g_4 can be viewed as parallel specifications, their reliabilities can be combined as follows:

$$R_{3,4} = 1 - (1 - R_3)(1 - R_4) \tag{18}$$

$$= 1$$

For the total reliability of P_1, consider R_1, R_2, and $R_{3,4}$ in series:

$$R_{tot,1} = R_1 R_2 R_{3,4}$$

$$= 0.8413$$

Calculation of the RSI of P_1 follows directly from this:

$$I_1(R) = -log(R_{tot,1}) \tag{19}$$

$$= 0.0750$$

Similar analysis can be used to evaluate the RSI of P_2 and P_3. Calculation of the RSI yields the following results: For P_2:

$$R_1 \simeq 1; R_2 \simeq 1; R_{3,4} = 0.75$$

$$I_2(R) = 0.1249$$

For P_3:

$$R_1 \simeq 1; R_2 = 0.1587; R_{3,4} \simeq 1$$

$$I_3(R) = 0.7994$$

Therefore, from a reliability standpoint, P_1 should be selected. The results are summarized in Table 1.

Further analysis of these results shows that each plan results in a nonzero probability of violating one of the operating specifications while meeting the other two specifications with almost perfect certainty: P_1 will perform the mission and avoid a collision with nearly perfect reliability, but results in a 16% chance of exceeding the maximum velocity; P_2 has a 25% chance of colliding with the moving obstacle, but will meet the mission specification and stay within the velocity bounds with nearly perfect certainty; and P_3 results in an 84% chance of not meeting the mission specification, but will stay within the velocity bounds and avoid collision with nearly perfect reliability.

Since no cost penalties are included in the reliability analysis, and since each plan results in a nonzero probability of failure of only one specification, P_1 is chosen because it offers a higher reliability with respect to the velocity constraint than does either P_2 with respect to the collision specification or P_3 with respect to the required task.

It is clear from this analysis that P_1 offers the plan with the highest probability of meeting all task constraints. However, looking at the costs associated with violation of the specifications as stated in the problem statement may complicate this result. Although P_1 offers the greatest probability of meeting all three specifications, the constraint it has the highest probability of violating is the velocity constraint. The assigned cost values show that this specification has the highest associated cost. In this scenario, it may be preferable to select another plan; one which is not as reliable, but has a higher probability of meeting the costly velocity constraint, while relaxing the less costly collision or mission specifications. For this type of analysis, a decision based on SSI may be employed.

3.3 Safety Analysis

Safety analysis is performed using the methods described in Section 2.2. The analysis is as follows:

First, costs are normalized:

$$\begin{aligned} c_1 &= \frac{\$900}{\$200} \\ &= 4.5 \\ c_2 &= \frac{\$200}{\$200} \\ &= 1.0 \\ c_{3,4} &= \frac{\$500}{\$200} \\ &= 2.5 \end{aligned}$$

These cost values are now used for calculation of the safety index, Ψ. For P_1:

$$\Psi_1 = \frac{5 - \mu_v}{\sqrt{((-1)(1)(c_1))^2}} \tag{20}$$

With $\mu_v = 5.0$ and $c_1 = 4.5$, Ψ_1 can be computed as:

$$\Psi_1 = 0.4938$$

Evaluating the cdf yields a safety factor, S_1, of:

$$S_1 = \Phi(\Psi_1) \tag{21}$$
$$= 0.6983$$

Similar analysis can be used for g_2, g_3, and g_4. Computation yields the following results:

$$S_2 \simeq 1.0$$
$$S_3 = 0.0228$$
$$S_4 = 0.9772$$

As with the RSI calculation, reduction of parallel specifications and combination of series specifications can be used to yield the total safety factor of P_1:

$$S_{tot,1} = 0.6827$$

From this, the SSI can be found directly:

$$F_1(S) = -log(S_{tot,1}) \tag{22}$$
$$= .1658$$

Repeating this analysis for P_2 and P_3 yields:

$$F_2(S) = .1281$$
$$F_3(S) = .7994$$

These results are summarized in Table 2. Choosing the plan with the lowest associated SSI results in selection of P_2. Although it has been shown that P_1 is the most reliable plan, safety analysis shows that P_2 has the lowest associated risk; from this definition of safety, P_2 is the safest plan. Although it allows for a higher probability of violating the mission specification than does P_1 with the velocity specification, the lower cost of violating the mission specification outweighs the higher probability of violation. In effect, this analysis has tightened the bounds on the velocity constraint to account for its higher associated cost.

3.4 Consideration of the Theory of Intelligent Machines

As was stated in Section 2.3, the Theory of Intelligent Machines is a design and analysis tool developed by Saridis to provide a theoretical structure for intelligent control systems. Safety analysis using SSI was shown to be consistent with this theory. This can be further demonstrated by considering the example problem.

Combination of Equations 3 and 9 yield the following relationship:

$$\Psi_{i,k} = \frac{\beta_{i,k}}{c_k} \tag{23}$$

This implies that for specification k with associated cost c_k, the safety index Ψ of plan i is directly proportional to the reliability index β of plan i with respect to specification k. In the example case, the largest cost is associated with the velocity constraint; therefore, when evaluating each of the three plans, the velocity measurement is treated as most uncertain when evaluating the safety index associated with specification 1. As shown in Equation 23, the use of a high cost value results in a decreased safety index and a corresponding increase in the SSI, indicating a decreased level of safety. As suggested by the Principle of Increasing Precision with Decreasing Intelligence presented in Section 2.3, this decreased level of safety must be countered with an increase in control performance; in this case, Equation 23 shows that the selection of a plan with a high reliability index with respect to specification 1 will counter the decrease in the level of safety caused by the cost-induced uncertainty. Using this result, safety analysis using the SSI can be viewed as a method of selecting plans which yield the highest weighted combination of specification reliabilities, requiring more reliable control with regard to more costly specifications. As indicated in Tables 1 and 2, plan P_2 is judged as the safest plan since it is highly reliable with respect to costly constraint 1, even though its overall reliability is lower than that of P_1, which is less reliable with respect to constraint 1. By considering control reliability to be a measure of precision and cost-induced uncertainties as a decrease in intelligence, the Principle of Increasing Precision with Decreasing Intelligence can be seen to manifest itself in safety analysis using SSI; more reliable control performance is expected in response to greater cost-induced uncertainties.

4 Conclusion

This paper has presented a quantifiable approach to safety for autonomous systems. A review of reliability theory has been presented, and the augmentation of reliability theory with cost information has been proposed. The concept of safety self-information has been defined, and has been demonstrated in a decision-making structure for a mobile robot. The safety analysis based on safety self-information has been shown to be consistent with the principle of Increasing Precision with Decreasing Intelligence.

Research will be continuing in the development of a quantifiable approach to safety. In autonomous environments, data sampling is often used as a means of collecting information about an unstructured environment. Probability distributions determined from finite data sets contain a degree of uncertainty characterized by *confidence levels*; this uncertainty will be used to augment the safety analysis presented in this paper. Also, this analysis does not address the issues of safety problems encountered due to failure of hardware and software components of the system; these component reliabilities will also be included in the safety analysis. In addition, the current safety analysis can only analyze existing plans; future research may include the use of the SSI to formulate plans. Future research may also address the computation time issues involved in safety analysis; often quick decisions must be made which cannot allow for a full analysis. In these cases, the need to perform an analysis must be weighed against the urgency of the decision at hand.

Acknowledgments

This work was supported by NASA grants NAGW-1333 and NGT-50418.

References

[1] James H. Graham, "Research Issues in Robot Safety", *Proceedings of the 1988 IEEE International Conference on Robotics and Automation,*

pp. 1854-1855, April 1988.

[2] John E. McInroy and George N. Saridis, "Reliability Analysis in Intelligent Machines", *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 20, No. 4, pp. 950-956, July/August 1990.

[3] John E. McInroy and George N. Saridis, "Reliable Control and Sensor Fusion in Intelligent Machines", *Proceedings of the 1991 IEEE International Conference on Robotics and Automation*, pp. 487 - 492, April 1991.

[4] J.E. McInroy and G.N. Saridis, "Reliability Analysis in Intelligent Machines", CIRSSE Document #39, Rensselaer Polytechnic Institute, August 1989.

[5] George N. Saridis, "On the Theory of Intelligent Machines: A Survey", *Proceedings of the 27th Conference on Decision and Control*, pp. 1799-1804, December 1988.

[6] G.N. Saridis, "An Integrated Theory of Intelligent Machines by Expressing the Control Performance as Entropy", *Control Theory and Advanced Technology*, Vol. 1, No. 2, pp. 125-138, August 1985.

[7] Alfredo H-S. Ang and Wilson H. Tang, *Probability Concepts in Engineering Planning and Design: Volume II - Decision, Risk, and Reliability*, John Wiley and Sons, New York, 1984.

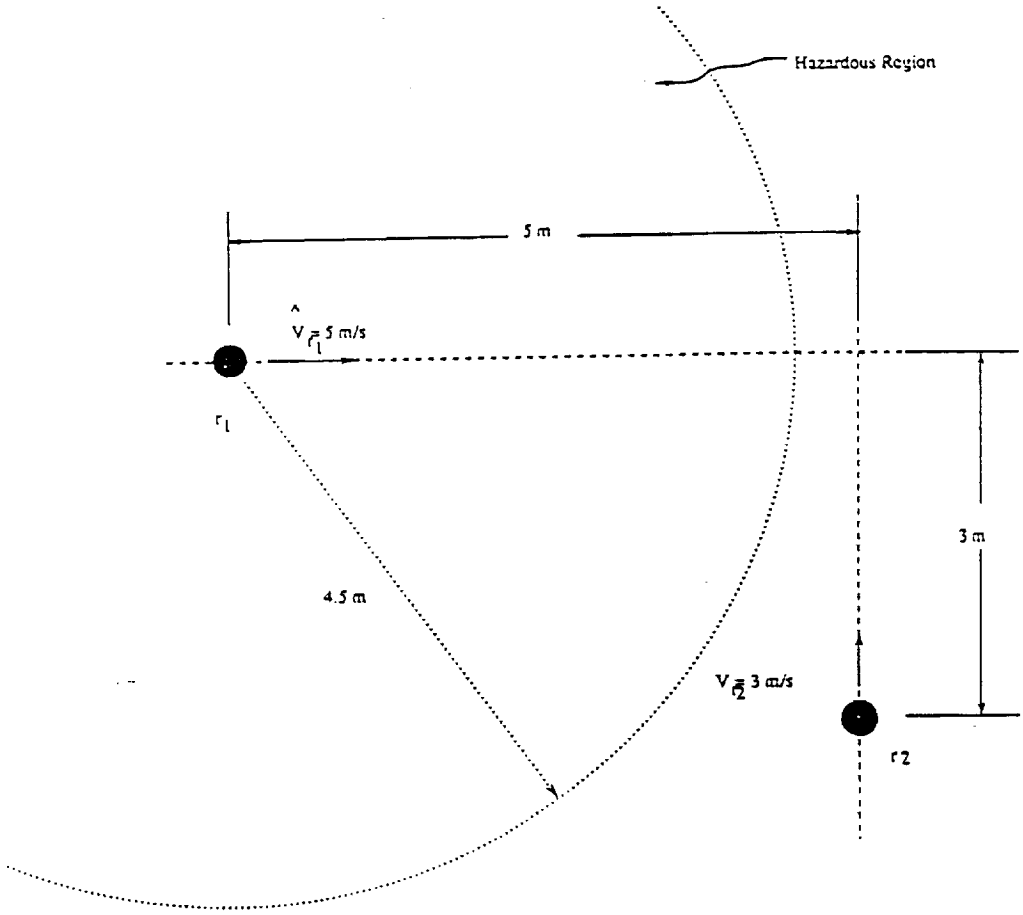

Figure 1: Mobile Robot Scenario

Table 1: Reliability Analysis

Plan	R_1	R_2	R_3	R_4	R_{tot}	$I(R_{tot})$
P_1	0.8413	1.0000	0.0000	1.0000	0.8413	0.0750
P_2	1.0000	1.0000	0.5000	0.5000	0.7500	0.1249
P_3	1.0000	0.1587	1.0000	0.0000	0.1587	0.7994

Table 2: Safety Analysis

Plan	S_1	S_2	S_3	S_4	S_{tot}	$F(S_{tot})$
P_1	0.6983	1.0000	0.0228	0.9772	0.6827	0.1658
P_2	0.9927	1.0000	0.5000	0.5000	0.7445	0.1281
P_3	1.0000	0.1587	1.0000	0.0000	0.1587	0.7994